

Expressing social attitudes in virtual agents for social training games

Nicolas Sabouret
LIMSI-CNRS, Univ. Paris-Sud
nicolas.sabouret@limsi.fr

Hazaël Jones
Univ. Pierre & Marie Curie
hazael.jones@lip6.fr

Magalie Ochs
Telecom ParisTech
magalie.ochs@telecom-paristech.fr

Mathieu Chollet
Telecom ParisTech
mathieu.chollet@telecom-paristech.fr

Catherine Pelachaud
Telecom ParisTech
catherine.pelachaud@telecom-paristech.fr

ABSTRACT

The use of virtual agents in social coaching has increased rapidly in the last decade. In order to train the user in different situations than can occur in real life, the virtual agent should be able to express different social attitudes. In this paper, we propose a model of social attitudes that enables a virtual agent to reason on the appropriate social attitude to express during the interaction with a user given the course of the interaction, but also the emotions, mood and personality of the agent. Moreover, the model enables the virtual agent to display its social attitude through its non-verbal behaviour. The proposed model has been developed in the context of job interview simulation. The methodology used to develop such a model combined a theoretical and an empirical approach. Indeed, the model is based both on the literature in Human and Social Sciences on social attitudes but also on the analysis of an audiovisual corpus of job interviews and on post-hoc interviews with the recruiters on their expressed attitudes during the job interview.

Keywords

Social Attitudes, Emotions, Affective computing, Virtual Agent, Non-verbal behaviour

1. INTRODUCTION

Social coaching workshops, organized by inclusion associations, social enterprises or coaching companies, constitute a common approach to help people in acquiring and improving their social competencies. However, it is an expensive and time-consuming approach that relies on the availability of trained practitioners as well as the willingness of the people to engage in exploring their social strengths and weaknesses in front of their peers and practitioners.

The use of virtual agents in social coaching has increased rapidly in the last decade. Projects such as those by Tartaro

and Cassell [29] or e-Circus [3] provide evidence that virtual agents can help humans improve their social skills and, more generally, their emotional intelligence [13].

Most of these models of virtual agent in social coaching domain have focused on the simulation of emotions [29, 3, 14]. However, in social coaching domain, a virtual agent may embody different social roles, such as virtual recruiter in a job interview simulation. Given its role and the course of the interaction, the virtual agent should be able to express different social attitudes to train the user in different situations that can occur in real life. Thus, one of the key elements of a virtual agent in the domain of social coaching is then its capacity to reason on social attitudes and to express them through its behaviour. This is why, in this paper, we propose a model of social attitudes for expressive virtual agents.

Social attitude can be defined as “*an affective style that spontaneously develops or is strategically employed in the interaction with a person or a group of persons, colouring the interpersonal exchange in that situation (e.g. being polite, distant, cold, warm, supportive, contemptuous)*” [23]. As highlighted in [26, 32], one’s social attitude depends on one’s personality but also one’s moods that is directly influenced by the course of the interaction. One’s social attitude is mainly conveyed by one’s non-verbal behaviour [2]. For instance, a dominant attitude is expressed through large and frequent gestures, a head directed upwards and larger amounts directed towards the interlocutor.

In this paper, we propose a model of social attitudes that enables a virtual agent to adapt its social attitude during the interaction with a user. It considers user’s emotions, mood and personality and compute and display agent’s appropriate social attitudes. Moreover, the model enables the virtual agent to display its social attitude through its non-verbal behaviour. The proposed model has been developed in the context of job interview (JI) simulation. This context is interesting for several reasons: 1) social attitude plays a key role in JI: the applicant tries to fit the social norm and 2) the recruiter is in a role-play game, which can be simulated with a virtual agent. Moreover, job interview is a type of social coaching situation with high social impact. The methodology used to develop such a model combined a theoretical and an empirical approach. Indeed, the model is based both on the literature in Human and Social Sciences on social attitudes but also on the analysis of an audiovisual

corpus of job interviews and on post-hoc interviews with the recruiters on their expressed attitudes during the interview.

The paper is organized as follows. In the next section, we discuss existing works in affective computing and intelligent virtual agent for social coaching. Section 3 presents the general architecture including a cognitive module to reason on social attitudes and an expressive module for the display. Section 4 presents the corpus that was collected to design the model. The cognitive module to compute the social attitude to express is introduced Section 5. The expressive module to give the capability to the virtual agent to adapt its non-verbal behaviour to convey social attitude is presented Section 6. We conclude Section 7.

2. RELATED WORK AND THEORETICAL BACKGROUND

A social attitude can be defined as “an affective style that spontaneously develops or is strategically employed in the interaction with a person” [23]. Social attitudes can be described by categories, but also can be represented along 2 dimensions: friendliness and dominance as in the interpersonal circumplex introduced by Leary [18]. Several research has shown that one’s social attitude is influenced by one’s affective state (e.g. his emotions and moods [11]) and the course of the interaction (e.g. the affective reaction of the other [32]). For instance in [32], Wegener et al. show that a positive mood can help influence a change of attitude in the interlocutor, and that people tend to feel a higher likelihood toward interlocutors that are in a positive mood.

Our aim is to develop a model of social attitudes that can be used by virtual agents to select determine expressive behaviours given its simulated affective state and the course of the interaction. Several models have been proposed in the domain of affective computing to compute the virtual agent’s emotions, moods and social relations considering different personality. More recently, several applications of intelligent virtual agents and affective models in the context of job interviews have been proposed [14, 5, 1]. However, most of these models are based on a reactive approach that connect expressed attitudes to perceptions in a reactive manner. On the contrary, our model propose to build and reason on a representation of social attitudes, so as to develop strategic intentions in the context of job interview.

Several research in Human and Social Sciences has shown that most modalities of the body are involved when conveying attitudes: smiles can be signs of friendliness[6], performing large gestures may be a sign of dominance, and a head directed upwards can be interpreted with a dominant attitude [7]. However, an attitude is not solely displayed by a sign. It is important to consider the succession of signs displayed by the agent as well as the signs displayed by the interlocutors. it is also crucial to consider how the signs of both interlocutors relate to each others. For example, it is only by looking at the sequencing of smile, gaze and head aversion that we can differentiate between amusement, shame and embarrassment, affects expressing different values of dominance [17].

Models of social attitude expression for virtual agents have already been proposed. For instance, in [4], postures cor-

responding to a given attitude were automatically generated for a dyad of agents. Ravenet *et al.* [22] proposed a user-created corpus-based methodology for choosing the behaviours of an agent conveying an attitude along with a communicative intention. The SEMAINE project used ECAs capable of mimicking and reacting to the system user’s behaviour to convey a personality. Each of these works used either limited modalities, or a non-interactive (agent-agent) context. Also, none of these works looked at the sequencing of the agent’s signal. In this article, we present model of social attitude expression that considers the sequencing of non-verbal behaviour (Section 6).

3. GENERAL ARCHITECTURE

Our approach of social coaching is based on role playing games (or interview simulation) with post-interview coaching sessions with an expert. Our goal is to provide the users with a simulation platform that allow people to train against a virtual agent. The coaching aspect is not part of this paper: we focus on the simulation platform and more precisely on the human-virtual agent interaction. Our goal is to define an architecture that supports intelligent expression of social attitudes, based on a cognitive architecture and a set of non-verbal behaviour bricks.

In this model, the interview simulation for social coaching involves two main actors, namely the participant (*i.e.* the person that is training on the system) and the interlocutor (*i.e.* person or virtual agent that respond to the trainee). In our platform, the interlocutor is replaced by a virtual agent.

Although the model presented here is general and can be applied to different interaction situations, our corpus and the derived cognitive architecture and non-verbal behaviour are designed in the context of job interview simulations when the agent acts as recruiter.

Fig. 1 presents our general architecture.

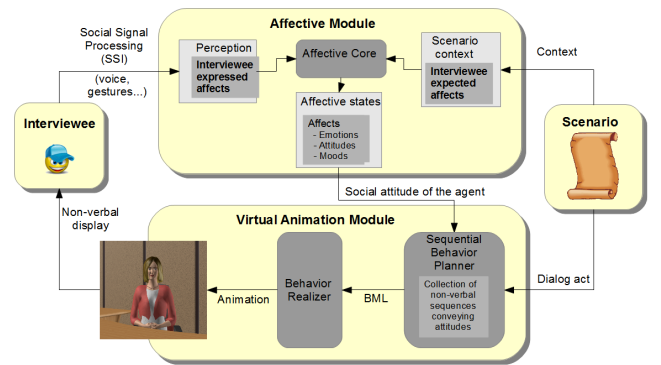


Figure 1: General architecture

This architecture is organized as follows. First, a *scenario* describes the context of the interaction and the expectations at the different stages of the interview. For example, after a tough question, the interviewer will expect negative affects from the interviewee (distress, agitation). These elements are used by the *Affective Module*. The interaction loop starts from the recognition of affects expressed by the

interviewee. Indeed, during a job interview, the emotions expressed by the interviewee constitute a key element influencing the recruiter’s attitudes, the interviewee’s emotions reflecting his social skills. In [25], a study shows that people who tried to suppress or hide negative emotions during a job interview are considered more competent by evaluators. Consequently, the *Affective Module* takes as input the affective state of the interviewee. Several tools have been developed to obtain this information in real-time. In our work, we use the SSI framework [31] that recognizes the user’s emotions expressed through his voice and his facial expressions.

The *Affective Module* (Section 5) reasons both on the affects (of the user and the virtual agent) and on the context of the interaction to derive the virtual agent’s social attitudes. These attitudes are turned into non-verbal behaviours in the virtual agent *Animation Module* (Section 6) and used in the next interaction step. The Animation Module is based on the SEMAINE platform [24], which includes an international common multi-modal behaviour generation framework. The *Intent Planner* generates the communicative intentions (what the agent intends to communicate). This information is provided by the scenario. The social attitude to express is transmitted to the *Behaviour Planner*. The *Behaviour Planner* transforms these communicative intentions into a set of signals (e.g. speech, gestures, facial expressions). Moreover, the *Behaviour Planner* has been extended to give the capability to the agent to display different social attitudes (Section 6). Finally, the *Behaviour Realizer* outputs for each of these signals the animation parameters.

In the following sections, we first present the corpus that was used to build our social attitude model. We explain how the corpus was collected, based on mock job-interviews, and it was annotated. We then give details of the Affective Module and the Animation Module and we show their interrelation.

4. CORPUS

We have collected a corpus of real interpersonal interaction between professional recruiters and job seekers interviewees. The recordings consisted in creating a situation of job interviews between 5 recruiters and 9 interviewees. The setting was the same in all videos. The recruiter and the interviewee sat on each side of a table. A single camera embracing the whole scene recorded the dyad from the side. This resulted in a corpus of 9 videos of job interview lasting between 15 and 20 minutes each. We discarded 4 videos as the recruiter was not visible due to bad position of the camera. Out of the 5 remaining videos, we have so far annotated 3, for a total of 50 minutes and 10 seconds of video.

The non-verbal behaviours of the recruiters during the job interview have been annotated. We also annotate information on the interaction: the *turn taking* (i.e. who is speaking), the *topic of the discussion* (i.e. document related or general), the perceived affects of the interviewee (e.g. embarrassed, relieved, bored...), and the attitude of the recruiter (i.e. the level of perceived dominance and friendliness of the recruiter). Different modalities of the recruiter’s non-verbal behaviour have been annotated (e.g. the gaze behaviour, the gesture, the posture, the head movements, etc). The coding

scheme, the resulting annotations, and the inter-annotators agreements are described in more details in [8]. These annotations have been used to construct the animation module for the virtual agent’s expression of social attitudes (Section 6). Moreover, the annotation on the recruiter’s social attitude has been used for the evaluation of the affective model to check that it produces outputs that correspond to the human behaviour (Section 6).

In addition to these videos, post-hoc interviews with the recruiters were used to elicit knowledge about the expectations and mental states during the interview, following the methodology proposed by [21]. This knowledge was used in the affective module to select the relevant social attitudes and to set up the rules to give the capability to the virtual agent to select the appropriate social attitude to express.

5. REASONING ABOUT ATTITUDES

5.1 Overview of the model

The Affective Module is based on a set of rules that compute categories of emotions, moods and attitudes for the virtual recruiter, based on the contextual information given by the scenario and the detected affects (emotions, moods and attitudes) of the participant (Section 3). The computation of the virtual agent’s emotions is based on the OCC model [20] and the computation of the agent’s moods is based on the ALMA model [12]. The details of the computation of emotions and moods will not be presented in this paper; it can be found in [16].

Formally, all affects in our model correspond to variables taking values in the interval $[0, 1]$ and we denote \mathcal{A} the set of all affects (moods, emotions and attitudes), with \mathcal{A}^+ the set of positive affects (joy, focused, etc) and \mathcal{A}^- the set of negative ones (distress, anxious, etc). Let $\mathcal{A}_d(t) \subset \mathcal{A}$ be the set of affects that are detected (d) at a time t of the interaction. We separate them in three subsets $E_d(t)$ (emotions), $M_d(t)$ (moods) and $A_d(t)$ (attitudes) and we denote $val(a)$ the value associated to each detected affect a . Similarly, we receive from the scenario a set of expected (e) affects $\mathcal{A}_e(t) \subset \mathcal{A}$. Expected affects are the expectations of the interviewer about the interviewee in term of affects, they are directly related to the nature of the question. An easy question will lead to positive expected affects and a tough question to negative ones. The values of the virtual recruiter’s simulated affects, $\mathcal{A}_f(t)$ (“f” stands for felt), is computed using expert rules based on the values of $\mathcal{A}_d(t)$ and $\mathcal{A}_e(t)$. All these rules are described in [16].

The list of affects (emotions, moods and attitudes) that are represented in the model for the virtual recruiter is given in Table 1. Note that this set is different from the detected and expected affects. It is based on the literature and on the mock interview corpus analysis (especially the knowledge elicitation phases presented in section 3). The emotions are a simple subset of the OCC model [20] that was selected based on what practitioners expressed, to different degrees, during the mock interviews. The moods originated from the ALMA model [12], limited to the positive dominance zone (since recruiters do not show submissive moods in the context of job interviews) with an additional distinction between bored and disdainful, that correspond to two subspaces of (P-,A-,D+) zone that were distinguished during

the knowledge elicitation phase. The attitudes were identified during the same knowledge elicitation phase and will be discussed in subsection 5.1.1.

	Positive	Negative
Emotions	Joy	Distress
	Relief	Disappointment
	Admiration	Anger
	Hope	Fear
Moods	Relaxed	Hostile
	Exuberant	Bored
		Disdainful
Attitudes	Friendly	Aggressive
	Supportive	
		Dominant
	Attentive	Inattentive
		Gossip

Table 1: Recruiter affects

The computation of moods is based on emotions following ALMA [12]. In the context of our interview simulation, the period is determined by the number of cycle question/answer. Each answer leads to the computation of a new emotions set and these emotions slightly influence the interviewer’s mood. The basis for calibration is as follows: after 5 cycles of a specific emotion (anger for example), the virtual recruiter will be in the corresponding mood (hostile). More details can be found in [16].

5.1.1 Virtual recruiter’s social attitudes

Several research has shown that one’s social attitude is influenced by one’s affective state (e.g. his emotions and moods [11]) and the course of the interaction (e.g. the affective reaction of the other [32]). For instance in [32], Wegener et al. show that people tend to feel a higher likelihood toward interlocutors that are in a positive mood.

Relations between attitudes and personality [26] and moods and attitudes [32, 11] has been exhibit in literature. Although we cannot give all the details of these papers here, their results show that one’s mood has an influence not only on the interlocutor’s attitude, but also on one’s own reaction to events. This knowledge has been turned into expert rules that compute values for the attitude of the virtual agent (using the categories elicited in the corpus and presented on table 1). These rules for the computation of attitudes rely on two main parameters: the actual mood of the virtual recruiter (which evolves during the simulation according to agent’s emotions) and the initial personality of the virtual recruiter (which will remain in the same state during the simulation). They are presented hereafter.

For the computation of the influence of personality on attitude, based on [26], we propose to use the five model [9] to represent the virtual recruiter’s personality using 5 quantitative dimensions (Openness (O), Conscientiousness (C), Extraversion (E), Agreeableness (A) and Neuroticism (N)).

Computation of attitudes

The way we compute attitudes follow this principle: an agent can adopt an attitude according to its personality [26]

or according to its actual mood [32]. For example, an agent with a non-aggressive personality may still show an aggressive attitude if its mood becomes very hostile. The mood compensates the personality and vice versa. For this reason, we use a logical-OR as condition on these two dimensions. As a consequence, in our model, the attitude can be triggered by one of these two dimensions. Then, the maximum value (mood or personality) is kept to compute the corresponding attitude, as is classically done in Fuzzy logics.

The minimum value for a trait to have an influence on the attitude is represented by a simple threshold θ . In our experiments, this threshold is a constant value set to 0.5; but it could be a function or a value that evolves in time. The following paragraphs present the computation rules for each attitude.

According to [9], friendly attitudes are more present in agreeable persons. Moreover, according to [15], positive and aroused moods are associated to the friendly quadrant. For these reasons, we define $val(friendly_f)$ as the combination of $val(A)$ and $val(exub.f)$. If $(val(A) > \theta) \vee (val(exub.f) > \theta)$, then $val(friendly_f) = \max(val(exub.f), val(A))$.

According to [30], aggressive attitudes are more present in non-agreeable and neuroticistic personalities and is correlated to the *hostile* mood. For this reason, if $((val(A) < \theta) \wedge val(N) > \theta) \vee (val(hostile_f) > \theta)$, then $val(aggr.f) = \max(val(hostile_f), val(N), 1 - val(A))$.

Similarly, we define dominant attitudes as: if $((val(E) > \theta) \wedge val(N) > \theta) \vee (val(hostile_f) > \theta)$, then $val(dominant_f) = \max(val(hostile_f), val(N), val(E))$.

For the same reasons, we define supportive attitudes as: if $((val(E) > \theta) \wedge val(A) > \theta) \vee (val(relax_f) > \theta)$, then $val(supportive_f) = \max(val(relax_f), val(A), val(E))$.

The attentive and inattentive dominance are simply computed based on the *conscientiousness* and the *disdainful* and *relaxed* mood, based on the definition of these affects:

- If $(val(C) < \theta) \vee (val(disd.f) > \theta)$, then $val(inatt.f) = \max(val(disd.f), 1 - val(C))$
- If $(val(C) > \theta) \vee (val(relax_f) > \theta)$, then $val(att.f) = \max(val(relax_f), val(C))$

Last, $val(gossip_f)$ was defined using *extraversion* and the *exuberant* attitude. If $(val(E) > \theta) \vee (val(exub.f) > \theta)$, then $val(gossip_f) = \max(val(exub.f), val(E))$.

5.2 Attitude on interpersonal circumplex

The non-verbal behaviour model of our agent, presented in the next section, does not work directly with the categories that were identified in the Knowledge elicitation phase of the corpus collection. On the contrary, it makes use of continuous values; it relies on the annotation of corpus which uses the Friendly and Dominant dimensions of the interpersonal circumplex [18].

To convert the attitudes represented by categories into continuous values of dimensions, we rely on the work by Isbister [15], summarized in Fig. 2. Our set of attitudes is shown in red on the circumplex, on the line from the center of the

circumplex to the outer. The position on this line relies on the intensity of the attitude.




Figure 2: Interpersonal circumplex by Isbister

When several attitudes are triggered at the same time, we compute the global attitude (that is the attitude that emerge) as the average of the associated vectors of the various attitudes. The vectors magnitudes influence this average giving more importance to an attitude with a large magnitude (i.e., intensity)). Considering n attitudes a_i :

$$Friendly = \frac{1}{n} \sum (val(a_i).Attitude_{Friendly}(a_i))$$

$$Dominance = \frac{1}{n} \sum (val(a_i).Attitude_{Dominance}(a_i))$$

As a result, the affective model gives a level of dominance and friendliness that represents the global attitude of the agent.

5.3 Evaluation of the affective model

To evaluate the affective module, we compare the affects that are computed from the affective module with the manual annotation of emotions from the job interview corpus. To perform such an evaluation, we chose arbitrarily one video of this corpus. The affective states of the interviewee and the social attitudes of the recruiter have been manually annotated. Our evaluation consists in comparing if, given the affective states of the interviewee, does the affective module computes social attitudes for the recruiter that are identical (i.e. similar in terms of dimensions representation) as the manually annotated social attitude of the recruiter in the video?

Let us describe briefly the sequence of 8 speaker turns in the video we are using for the evaluation: in the first 4 questions/answers turns, the interviewee is confident and gives good answers to the recruiter (positive detected affects). The interviewer expectations has been annotated positively during this first sequence. Then, the 4 following questions

appear more difficult for the interviewee; he shows expression of negative affects. On the other hand, expectations of the recruiter were annotated positive for questions 5 and 6 and negative for questions 7 and 8.

Fig. 3 and Fig. 4 show the recruiter affects. Fig. 3 shows the output for the recruiter's affects after each of its question. We can notice that positive affects (triangles) are positively correlated with supportive attitude while negative affects (squares) are positively correlated with aggressive attitude. Fig. 4 displays the annotated social attitudes of the recruiter as it evolves during the interaction. Positive values mean that the recruiter is perceived supportive by the user while negative values as aggressive. We also indicate where the 8 questions happen in the course of the interaction.

Comparing both data, the outputs of the affective module and the manual annotation, we can remark that the results are not really comparable in term of intensity of social attitudes. However the attitudes computed by the affective module coincide with the manually annotated attitudes. The variation of attitudes from supportive to aggressive happens in both cases. This example shows that our affective module computes, for a given input (the affect states of the interviewee), similar attitudes for the recruiter of those that are perceived in real human-human job interview.




Figure 3: Recruiter affects after each question




Figure 4: Recruiter state in real time

6. EXPRESSION OF ATTITUDES

Research has shown that non-verbal behaviours convey interpersonal attitude (see Section 2. However it is insufficient to look at signals independently of the other surrounding signals: a smile is a sign of friendliness, but a smile preceded

by head and gaze aversion conveys submissiveness [17]. It is also important to consider how signals from the virtual agent and signals from the user are sequenced (*e.g.* for mimicry).

In our work, to give the capability to a virtual agent to convey attitudes, we choose a corpus-based approach to find how attitudes are expressed through signal sequences. Rather than finding direct mapping of attitudes to sequences of behaviors, we are interested in characterizing when a sequence of behaviours changes the perception of the associated attitudes. It allows us to capture how attitudes evolve through the interaction through the dynamic display of nonverbal behaviours. Using sequence mining techniques, we extract sequences and characterize how likely they are to convey an attitude variation. The result of this analysis is integrated in the *Behaviour Planner* of the SAIBA architecture, that receives the attitude computed by the *Affective Module*, and chooses a sequence of non-verbal behaviours accordingly. In the following section, we present an algorithm that extract sequences of non-verbal signal.

6.1 Extraction of sequences from the corpus

In order to extract significant sequences of non-verbal signals conveying interpersonal attitudes from our corpus, we use a *frequent sequence mining* technique [27]. Such techniques have been widely used in protein classification [10], and recent work have used these techniques in multimodal contexts: in a video game context, Martinez and Yannakakis studied the sequences of user keystrokes, game related events and user physiological signals [19]. To the best of our knowledge, this technique has not yet been applied to analyse sequences of non-verbal signals. In order to apply the frequent sequence mining technique to our data, we proceed through the following six steps.

The first step consists of parsing the non-verbal annotations files (coded in the ELAN format), filtering the annotation modalities to investigate and converting every interaction’s annotations into a list containing all the non-verbal behaviours in a sequence.

The second step as for objective to find events to segment the interactions: indeed, frequent sequence mining techniques require a dataset of sequences. In our case, our data consists of 3 continuous interactions (videos of 3 full job interview sessions). Since we investigate which sequences of behaviour convey attitudes, we decide to segment the full interactions with attitude variation events: *attitude variation events* are the timestamps where the attitude begins to vary. To this end, we parse the attitude annotations files, smooth them and find the timestamps where an attitude starts to vary. More details can be found in [8].

Thirdly, we aggregate the attitude variation events into different clusters. Using a K-Means clustering technique, with $k = 4$. For friendliness and dominance, we obtain 4 clusters separating the attitude variation events in clusters for large increases, small increases, small decreases or large decreases in value.

The fourth step consists of segmenting the full interaction sequences with the attitude variation events obtained from step 2. Following this procedure, we obtain 245 segments

preceding dominance variations and 272 preceding friendliness variations. These two sets are split further depending on which cluster the attitude variation event belongs to. For instance, we have 79 segments leading to a large drop in friendliness, and 45 segments leading to a big increase in friendliness.

Step 5 consists of applying the frequent sequence mining algorithm to each set of segments. We used the commonly used Generalized Sequence Pattern (GSP) frequent sequence mining algorithm described in [27]. The output of the GSP algorithm is a set of sequences along with their support, *i.e.* the number of times a sequence happens in the dataset.

However, the support is an insufficient measure to analyse how a sequence is characteristic of an attitude variation: for instance, having the gaze move back and forth to the interlocutor happens very regularly in an interaction. Thus it will happen very often before an attitude variation (*i.e.* it will have a high support), even though it is not sure that it is characteristic of an attitude variation. The objective of step 6 is to compute *quality measures* to assess whether a sequence is really characteristic of an attitude variation. Based on [28], we choose to compute *confidence* and *lift* quality measures for every sequence. The confidence represents how frequently a sequence is found before an attitude variation, and the lift represents how the sequence occurs before an attitude variation more than expected (the higher the value, the more likely it is that there is dependence between the sequence and the attitude variation).

In Tab. 2 we show examples of extracted sequences found using this process, and the corresponding measures. The *Sup* column corresponds to the support of the sequence and the *Conf* column to the confidence of the sequence.

Sequence	Variation type	<i>Sup</i>	<i>Conf</i>	<i>Lift</i>
EyesAt-> EyesAway-> EyebrowUp->HeadShake	Friendliness Big Decrease	0.027	0.727	2.72
HeadAt->GestComm-> EyebrowUp->Smile	Dominance Big Increase	0.023	0.6	5.38

Table 2: Examples of extracted sequences

In the next section, we present how appropriate sequences are chosen using *lift* and *confidence* values.

6.2 Sequences selection for attitude expression

The attitude expression is handled in our virtual agent platform in the *Behaviour Planner* module. The original role of this module is to instantiate communicative intentions into different signals, such as gestures or facial expressions. In addition to this, we extend the module to have it plan sequence of signals that will convey a chosen attitude. As shown in Fig. 1, it receives attitude values computed by the *Affective Module*, in the form of a pair of values, *Friendly* and *Dominance*.

The *Behaviour Planner* computes $\Delta_{Friendly}(a_n, a_{n-1})$ and $\Delta_{Dominance}(a_n, a_{n-1})$, the differences between the chosen attitude in the last dialogue turn a_{n-1} and the new atti-

tude a_n to express. The *Behaviour Planner* then selects the set of sequences corresponding to the attitude variation. Moreover, we look for a non-verbal sequence that will be feasible in the current context. For instance, if the virtual agent is already leaning backwards, the potential sequences containing a change of posture backwards will be discarded. To select the appropriate sequence, for every candidate sequence seq and for target attitude variations $\Delta_{Friendly}$ and $\Delta_{Dominance}$, the following values are computed considering its confidence $Conf$ and lift $Lift$ measures:

$$Val(seq) = Conf(seq, \Delta_{Friendly}) * Conf(seq, \Delta_{Dominance}) * Lift(seq, \Delta_{Friendly}) * Lift(seq, \Delta_{Dominance}) / \lambda$$

with λ being a coefficient used to weigh down sequences if they have been recently displayed, in order to avoid repetition of the same sequences. The highest scoring sequence is chosen by the *Behaviour Planner* and passed on to the *Behaviour Realizer*, in turn producing the animation of the virtual agent. An example of such an animation is shown in Fig. 5.




Figure 5: Example of the non-verbal sequence $HeadAt \rightarrow GestComm \rightarrow EyebrowUp \rightarrow Smile$ (from Tab. 2) expressed by the virtual agent.

Our next step, to be performed in the very near future, consists of conducting user perceptive tests to validate that the sequences computed by the *Behaviour Planner* convey the chosen attitude when expressed by the virtual agent.

7. CONCLUDING REMARKS

In the context of social coaching with virtual agents, this paper proposes a new architecture for expressive agents that can reason about and display social behaviours. Unlike classical reactive models, our approach combines an affective reasoner that generates affects for the virtual character with a sequence selection mechanism based on a domain corpus annotated on two dimensions: dominance and friendliness.

The methodology we propose contains several stages, from corpus collection and annotation, knowledge elicitation with experts for the definition of rules, implementation of behaviours corresponding to sequences and sequence selection based on the generated internal affects. The originality of our model is twofold: first, we rely on an affective model that receives as inputs expectations on user affects, compares them to user perceptions and computes agent affects accordingly. All affects were designed after knowledge elicitation phases with domain experts. Second, the expression of social attitudes does not only consider the signals to convey such attitudes but also the non-verbal behaviours context. Contrary to most existing models of communicative behaviours for ECAs, in our model, to display an attitude, we generate a sequence of non-verbal signals taking into account the signals the agent has shown previously as well as the expressions displayed by its interlocutors. We used a corpus-based approach to identify sequences of non-verbal behaviours that display variations in attitudes.

This architecture, whose components have been tested separately, has been integrated using the SEMAINE platform [24] and is currently being tested with real users. This will allow us to validate the global behaviour of our platform in the context of social coaching. However, several components can still be improved. One first limit of our model is that we assume exact inputs from the perception module. Dealing with imprecision in social signal interpretation is one of our next steps and we intend to use probabilistic or fuzzy theories. In addition, we intend to provide the affective reasoner with a representation of the interaction from the recruiter's point of view. We believe that allowing the recruiter to reason about the actual and potential behaviour of the applicant, following a Theory of Mind paradigm, will allow a more credible decision process. We will also take the sequence extraction procedure further to take into account the user's non-verbal signals. This will require the *Behaviour Planner* sequence selection method to plan for the occurrence of user signals, and to react accordingly when the user does not react as planned.

8. REFERENCES

- [1] K. Anderson, E. André, T. Baur, S. Bernardini, M. Chollet, E. Chrysafidou, I. Damian, C. Ennis, A. Egges, P. Gebhard, H. Jones, M. Ochs, C. Pelachaud, K. Porayska-Pomsta, R. Paola, and N. Sabouret. The TARDIS framework: intelligent virtual agents for social coaching in job interviews. In *Proceedings of the Tenth International Conference on Advances in Computer Entertainment Technology (ACE-13)*. Enschede, the Netherlands. LNCS 8253, November 2013.
- [2] M. Argyle. *Bodily Communication*. University paperbacks. Methuen, 1988.
- [3] R. Aylett, A. Paiva, J. Dias, L. Hall, and S. Woods. Affective agents for education against bullying. In *Affective Information Processing*, pages 75–90. Springer, 2009.
- [4] D. Ballin, M. Gillies, and B. Crabtree. A framework for interpersonal attitude and non-verbal communication in improvisational visual media production. In *1st European Conference on Visual Media Production*, 2004.

- [5] L. M. Batrinca, G. Stratou, A. Shapiro, L.-P. Morency, and S. Scherer. Cicero - towards a multimodal virtual audience platform for public speaking training. In *Proc. 2013 International Conference on Intelligent Virtual Agents*, pages 116–128, 2013.
- [6] J. K. Burgoon, D. B. Buller, J. L. Hale, and M. A. de Turck. Relational Messages Associated with Nonverbal Behaviors. *Human Communication Research*, 10(3):351–378, 1984.
- [7] D. R. Carney, J. A. Hall, and L. S. LeBeau. Beliefs about the nonverbal expression of social power. *Journal of Nonverbal Behavior*, 29(2):105–123, 2005.
- [8] M. Chollet, M. Ochs, and C. Pelachaud. A multimodal corpus for the study of non-verbal behavior expressing interpersonal stances. In *IVA'13 Workshop Multimodal Corpora*, September 2013.
- [9] P. T. Costa and R. R. MacCrae. *Revised NEO Personality Inventory (NEO PI-R) and NEO Five-Factor Inventory (NEO FFI): Professional Manual*. Psychological Assessment Resources, 1992.
- [10] P. Ferreira and P. Azevedo. Protein sequence classification through relevant sequence mining and bayes classifiers. In C. Bento, A. Cardoso, and G. Dias, editors, *Progress in Artificial Intelligence*, volume 3808 of *Lecture Notes in Computer Science*, pages 236–247. Springer Berlin Heidelberg, 2005.
- [11] J. P. Forgas, G. H. Bower, and S. E. Krantz. The influence of mood on perceptions of social interactions. *Journal of Experimental Social Psychology*, 20(6):497–513, 1984.
- [12] P. Gebhard. ALMA - A Layered Model of Affect. *Artificial Intelligence*, pages 0–7, 2005.
- [13] D. Goleman. *Social intelligence: The new science of human relationships*. Bantam, 2006.
- [14] M. Hoque, M. Courgeon, J.-C. Martin, B. Mutlu, and R. Picard. MACH: My Automated Conversation coach. In *Proc. 2013 ACM international joint conference on Pervasive and ubiquitous computing*. ACM Press, 2013.
- [15] K. Isbister. *Better Game Characters by Design: A Psychological Approach (The Morgan Kaufmann Series in Interactive 3D Technology)*. Morgan Kaufmann Publishers Inc., 2006.
- [16] H. Jones and N. Sabouret. TARDIS - A simulation platform with an affective virtual recruiter for job interviews. In *IDGEI (Intelligent Digital Games for Empowerment and Inclusion)*, 2013.
- [17] D. Keltner. Signs of appeasement: Evidence for the distinct displays of embarrassment, amusement, and shame. *Journal of Personality and Social Psychology*, 68:441–454, 1995.
- [18] T. Leary. Interpersonal diagnosis of personality: a functional theory and methodology for personality evaluation. *Methodology*, 1957.
- [19] H. P. Martínez and G. N. Yannakakis. Mining multimodal sequential patterns: a case study on affect detection. In *Proceedings of the 13th international conference on multimodal interfaces*, ICMI '11, pages 3–10, New York, NY, USA, 2011. ACM.
- [20] A. Ortony, G. L. Clore, and A. Collins. *The Cognitive Structure of Emotions*. Cambridge University Press, July 1988.
- [21] K. Porayska-Pomsta, M. Mavrikis, S. D'Mello, C. Conati, and R. Baker. Knowledge elicitation methods for affect modelling in education. *International Journal of Artificial Intelligence in Education*, 2013.
- [22] B. Ravenet, M. Ochs, and C. Pelachaud. From a user-created corpus of virtual agent's non-verbal behaviour to a computational model of interpersonal attitudes. In *International Conference on Intelligent Virtual Agent (IVA2013)*, 2013.
- [23] K. R. Scherer. Appraisal considered as a process of multilevel sequential checking. *Appraisal processes in emotion: Theory, methods, research*, pages 92–120, 2001.
- [24] M. Schröder, E. Bevacqua, R. Cowie, F. Eyben, H. Gunes, D. Heylen, M. ter Maat, G. McKeown, S. Pammi, M. Pantic, C. Pelachaud, B. Schuller, E. de Sevin, M. F. Valstar, and M. Wöllmer. Building autonomous sensitive artificial listeners. *T. Affective Computing*, 3(2):165–183, 2012.
- [25] M. Sieverding. 'Be Cool!': Emotional costs of hiding feelings in a job interview. *International Journal of Selection and Assessment*, 17(4), 2009.
- [26] M. Snyder. The influence of individuals on situations: Implications for understanding the links between personality and social behavior. *Journal of Personality*, 51(3):497–516, 1983.
- [27] R. Srikant and R. Agrawal. Mining sequential patterns: Generalizations and performance improvements. In P. Apers, M. Bouzeghoub, and G. Gardarin, editors, *Advances in Database Technology (EDBT) 96*, volume 1057 of *Lecture Notes in Computer Science*, pages 1–17. Springer Berlin Heidelberg, 1996.
- [28] P.-N. Tan, M. Steinbach, and V. Kumar. *Introduction to Data Mining, (First Edition)*. Addison-Wesley Longman Publishing Co., Inc., Boston, MA, USA, 2005.
- [29] A. Tartaro and J. Cassell. Playing with virtual peers: bootstrapping contingent discourse in children with autism. In *Proceedings of the 8th international conference on International conference for the learning sciences-Volume 2*, pages 382–389. International Society of the Learning Sciences, 2008.
- [30] P. F. Tremblay and L. A. Ewart. The Buss and Perry Aggression Questionnaire and its relations to values, the Big Five, provoking hypothetical situations, alcohol consumption patterns, and alcohol expectancies. *Personality and Individual Differences*, 38(2):337–346, 2005.
- [31] J. Wagner, F. Lingenfelser, T. Baur, I. Damian, F. Kistler, and E. André. The social signal interpretation (ssi) framework-multimodal signal processing and recognition in real-time. In *Proceedings of the 21st ACM International Conference on Multimedia, Barcelona, Spain*, 2013.
- [32] D. T. Wegener, R. E. Petty, and D. J. Klein. Effects of mood on high elaboration attitude change: The mediating role of likelihood judgments. *European Journal of Social Psychology*, 24(1):25–43, 1994.